Development Of A Massively Parallel NOGAPS Forecast Model

Tom Rosmond
Naval Research Laboratory
Monterey, CA 93943-5502
phone: (831) 656-4736 fax: (831) 656-4769 e-mail: rosmond@nrlmry.navy.mil
Award #: N0001498WX30167

LONG-TERM GOAL

Develop an advanced global atmospheric forecast system designed to exploit massively parallel processor (MPP), distributed memory computer architectures. Future increases in computer power from MPP's will allow substantial increases in model resolution, more realistic physical processes, and more sophisticated data assimilation methods, all of which will improve operational numerical weather predictions and provide better simulations of the Earth's climate.

OBJECTIVES

The current Navy operational global atmospheric prediction system (NOGAPS 4.0) is a highly optimized Fortran code designed to run on parallel vector, shared memory machines (CRAY's). The immediate objective of the project is to redesign the model's numerical algorithms and data structures to allow efficient execution on MPP architectures and clusters of shared memory processors. This includes using icosahedral grids, finite element, spectral element, and Lagrangian methods. Message passing (MPI) is the paradigm chosen for communication between distributed memory processors.

APPROACH

Use integrations of the current operational NOGAPS as control runs to ensure reproducibility of results with the newly designed Fortran 90 code. Design efficient spectral transform algorithms for both shared memory and distributed memory architectures. For distributed memory architectures use message passing library modules in communication intensive spectral transforms and horizontal interpolation routines.

The current NOGAPS spectral formulation requires global communication for the spherical harmonic transforms. An attractive alternative is the use of quasi-uniform icosahedral grids based on local basis functions that are less communication intensive. Implementation of the icosahedral grid NOGAPS to distributed memory architecture and the addition of vertical coordinates to the current barotropic version have begun.

WORK COMPLETED

The scalable NOGAPS MPI code has been further refined and optimized to run on both distributed memory and shared memory architectures. A companion code using the recently established OpenMp programming paradigm for shared memory multi-processing has been developed, and work on merging the MPI and OpenMp codes into a single flexible NOGAPS code has begun. This is an important

Report Documentation Page					*Form Approved* *OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **30 SEP 1999**	2. REPORT TYPE	3. DATES COVERED **00-00-1999 to 00-00-1999**
4. TITLE AND SUBTITLE **Development Of A Massively Parallel NOGAPS Forecast Model**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Naval Research Laboratory,Monterey,CA,93943-5502**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**5**	

objective from a configuration management perspective to avoid the necessity of maintaining separate source codes for the two types of architectures, or to run NOGAPS on hybrid architectures such as clusters of SMP nodes that require both programming models.

The NOGAPS MPI code was provided to FNMOC to be used as the primary benchmark for their POPS upgrade procurement. Data sets, run-scripts, and user documentation were prepared and provided to potential vendors. The procurement process went quite smoothly, with almost no objections or complaints from any of the participating vendors about the NOGAPS benchmark package. Four vendors successfully ran the NOGAPS benchmark code, with Silicon Graphics being chosen as the winning vendor based on benchmark performance.

Spherical Harmonic transforms dominate the communication workload in the MPI NOGAPS. A new transform algorithm has been developed that scales to larger numbers of processors than the transform algorithm used in the benchmark version of the code. The new method was developed in collaboration with the group at the UCSD supercomputer center team support by the High Performance Computing PET program.

Shallow water versions of the finite element and spectral element icosahedral grid NOGAPS have been completed and extensively tested.

RESULTS

One of the objectives of the MPI NOGAPS code development was to be as faithful as possible to the current operational code run on the FNMOC Cray C90. Due to inherent differences in several computational algorithms between the MPI code and the shared memory Cray code, bit-reproducibility was not possible, but 8-9 significant figure matching between solutions from the two versions has been achieved after 12 hours of integration. This result is considered excellent, and demonstrated that the MPI NOGAPS is clearly equivalent to the operational model, and will be an excellent starting point for continued model development on the new FNMOC system when it becomes available in early FY2000.

Faithfully reproducing the operational NOGAPS code with the MPI NOGAPS code demonstrated a serious problem with one computational module of the model, however. The post-processing of model forecasts to produce standard FNMOC gridded fields requires horizontal and vertical cubic spline interpolations. The algorithms used are extremely efficient on shared memory vector architectures such as the C90, but have prohibitive MPI communication overhead on distributed memory scalable architectures. Fig. 1 shows model runtime for a 24-hour forecast, including the 2000 operationally produced fields. Notice that the cost of cubic spline interpolation ***increases*** dramatically with increasing processor numbers, clearly an unacceptable situation for operational implementation.

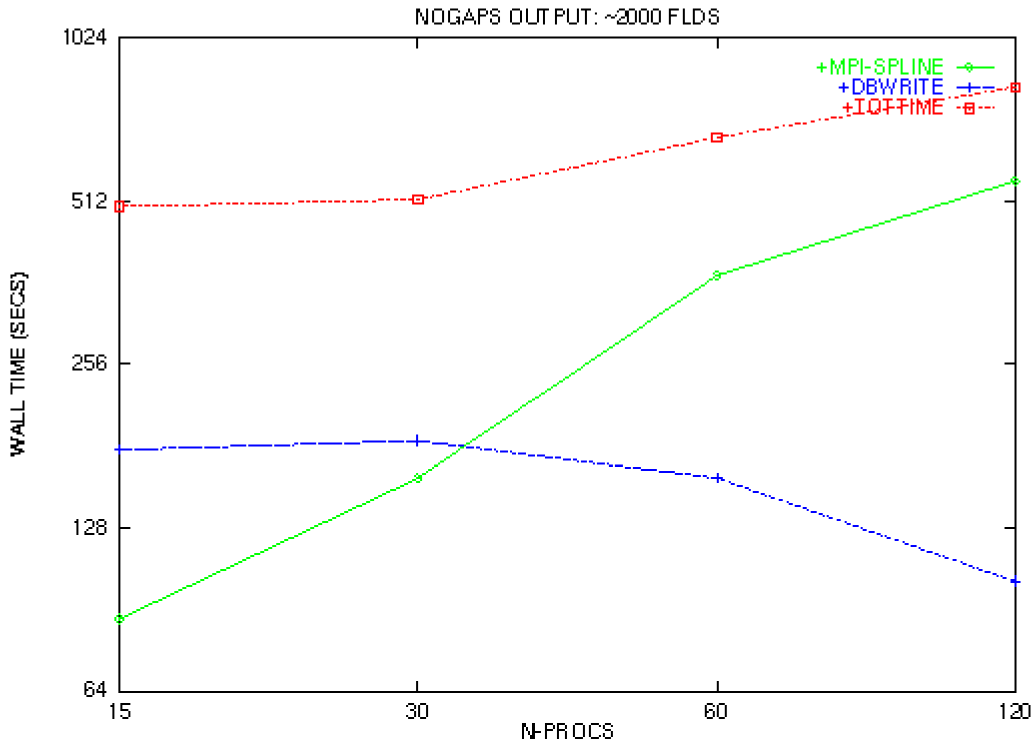

Figure 1: NOGAPS post processing time for 24 hours of output (2000 fields) on a Cray T3E. Shown is the total runtime, including time integration (tottime), IO time (dbwrite), and spline interpolation time (mpi-spline), which includes the communication cost. Note the dramatic increase in the latter as the number of processors increases, so that for 120 processors the splines take more time than the model integration. In comparison on 6 Cray C90 processors the same interpolations take about 90 seconds

There are at least two possible solutions to this problem. The cubic spline algorithm could be replaced with a less communication-intensive scheme, e.g. linear, but this would sacrifice accuracy and so is not attractive. A more likely solution is possible because the super-computing environment of the future will be more heterogeneous that today's. Scalable architectures are not general-purpose machines to nearly the same degree that machines such as the C90 have been, as the results here clearly demonstrate. The NOGAPS post-processing software is an ideal candidate for running off-line from the model time integration on stand-alone shared memory SMP systems, or even individual shared memory nodes of the scalable architecture that are not dedicated to the time integration. In either scenario, having a single source code that can run on either distributed memory or shared memory architectures makes this an operationally viable alternative, and is also attractive for NOGAPS research applications because of the flexibility it provides.

The icosahedral grid NOGAPS has been used to evaluate the relative merits of local finite element and local spectral element methods on these kinds of quasi-regular grids. A number of papers and presentations on the results have been published showing that this is a viable method for solving the atmospheric equations on the sphere accurately and efficiently. The figure below shows the spectral (exponential) convergence of the spectral element method. As the order of the polynomial "p" is increased by a factor of 2, the L2 error decreases by a factor of 10.

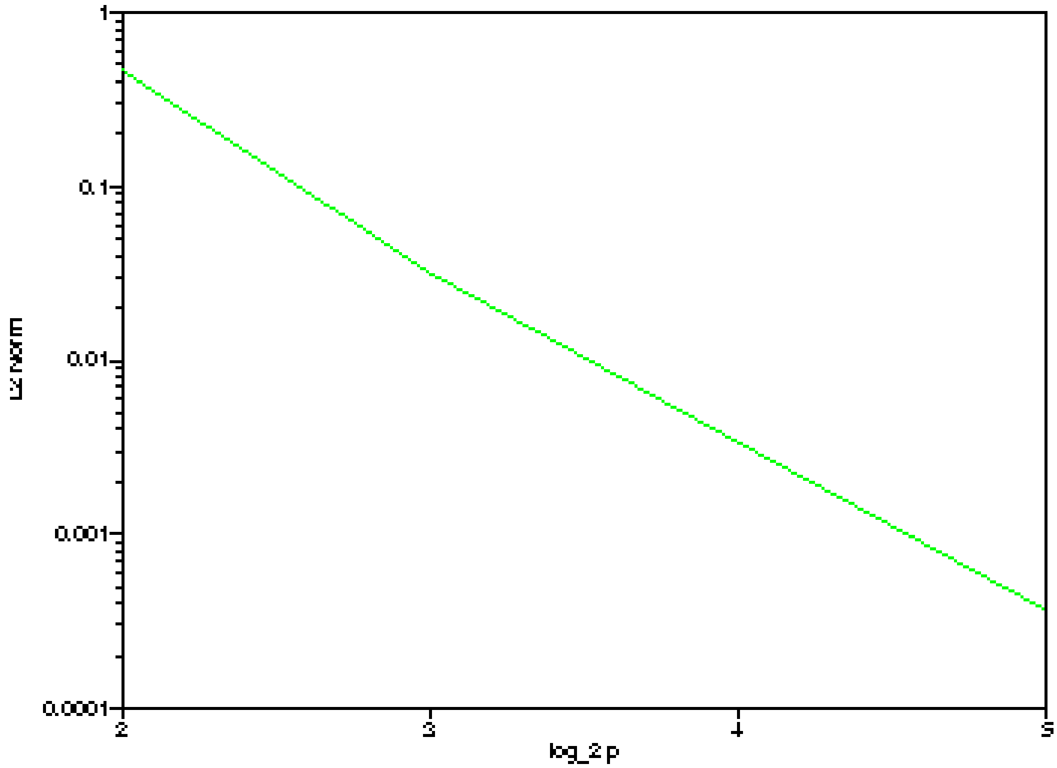

IMPACT

NOGAPS is run operationally by FNMOC and is the heart of the Navy's operational weather prediction support to nearly all DOD users worldwide. It is also run by many NRL and other Navy researchers to study atmospheric dynamics, and atmosphere/ocean interaction. Our work here targets the next generation of this system for the next generation of computer architectures. These architectures are expected to be distributed memory, commodity based systems with enormous theoretical computational power. However, exploiting this capability will require drastically redesigning many important model algorithms.

TRANSITIONS

Improved algorithms for model processes will be transitioned to 6.4 (PE 0603207N) as they are ready, and will ultimately be transitioned to FNMOC with future NOGAPS upgrades. Development of the MPI NOGAPS code has necessitated close examination of the algorithms used in the operational model, and in some cases uncovered design weaknesses and bugs that are being promptly corrected in the operational NOGAPS.

The scalable NOGAPS MPI code was provided to FNMOC for their POPS upgrade procurement. Scripts, initial data sets, and user documentation were also provided. The procurement went very smoothly, with virtually no complaints for vendors, and the selection of SGI as the winning vendor.

RELATED PROJECTS

(1) NOGAPS 4.0 Evaluation (X0513-01): Advanced development and transition of the NOGAPS 4.0 forecast model to operational status at FNMOC. (2) The DOD CHSSI Scaled Software algorithm development for meteorological models (HPCM-96-032): Development of numerical algorithms appropriate for massively parallel computer architectures. These algorithms will be critical for inter-processor communication dependent and computationally intensive model processes.

REFERENCES

Giraldo, F. X., 1998: Efficiency and Accuracy of Lagrange-Galerkin Methods on Unstructured Adaptive Grids. *Mathematical Modelling & Scientific Computing, Vol.* 8.

Giraldo, F. X., 1998: The Lagrange-Galerkin Spectral Element Method on Unstructured Quadrilateral Grids. *Journal of Computational Physics, Vol.* 147, pp. 114-146.

Giraldo, F. X., 1999: Trajectory Calculations for Spherical Geodesic Grids in Cartesian Space. *Monthly Weather Review, Vol.* 127, pp. 1651-1662.

Giraldo, F. X. and Neta, B., 1999: Stability Analysis for Eulerian and Semi-Lagrangian Finite Element Formulation of the Advection-Diffusion Equation. *Computers and Mathematics with Applications, Vol.* 38, pp. 97-112.

Giraldo, F. X., 1998: A Weak Lagrange-Galerkin Shallow Water Model on the Sphere. Presented at Los Alamos National Laboratory, Applied, Theoretical and Computational Physics Division (December 1998).

Giraldo, F. X., 1999: Lagrange-Galerkin Methods on Spherical Geodesic Grids: The Shallow Water Equations. Presented at Colorado State University, Department of Atmospheric Science (September 1999).

Rosmond, T. E., 1998: A Scalable Version of the Navy Operational Global Atmospheric Prediction System Spectral Forecast Model. Submitted to *Scientific Programming*. Oct. 1998.

Rosmond, T. E., 1998: Investigation of the Scalability of Pre-processing, Post processing, and Computational Modules of the NOGAPS Spectral Forecast Model. Presented at the ECMWF workshop: "Toward Teracomputing – The use of parallel processors in Meteorology". 16-20 November, 1998.